

Improving Mixed-Criticality System Consistency and Behavior on Multiprocessor Platforms by Means of Multi-Moded Approaches

François Santy

Geoffrey Nelissen

Joël Goossens

PARTS Research Center
Université Libre de Bruxelles (ULB)
Brussels, Belgium

Abstract—Recent research in the domain of real-time scheduling theory has tackled the problem of scheduling mixed-criticality systems upon uniprocessor or multiprocessor platforms, with the main objective being to respect the timeliness of the most critical tasks, at the expense of the requirements of the less critical ones. In particular, the less critical tasks are carelessly discarded when the computation demand of (some of) the high critical tasks increases. This might nevertheless result in system failure, as these less critical tasks could be accessing data, the consistency of which should be preserved. In this paper, we address this problem and propose a method to cautiously handle task suspension. Furthermore, it is usually assumed that the less critical tasks will never be re-enabled once discarded. In this paper, we also address this concern by proposing an approach to re-enable the less critical tasks, without jeopardizing the timeliness of the high critical ones. The suggested approaches apply to systems having two or more criticality levels.

I. INTRODUCTION

The current trend in embedded systems is towards collocating multiple functionalities on shared resources, as illustrated by industrial initiatives in both the aerospace and automotive industry. Nowadays, this trend is even getting stronger with the introduction of multiprocessor and/or multicore architectures. With such collocation however, it is unlikely that all functionalities share the same level of importance (*criticality*), and the timeliness of some functionalities might appear more important than others. In this context, mandatory certification of whole (or subset of) the system by statutory Certification Authorities (CAs) might be required to guarantee the correct behavior of the latter. When certifying a subset of the functionalities, the CAs have to make assumptions about the Worst-Case Execution Time (WCET) of tasks. In practice, it can be observed that the more a task will be assumed critical, the more will it be certified using strongly pessimistic assumptions about its behavior at run-time. In this context, *mixed-criticality* (MC-) systems are an attempt to model systems that need to be certified using *various* assurance levels. The increased pessimism assumed for the more critical tasks during the certification process however introduces a significant over-provisioning of resources at run-time for the latter, having as consequence that the scheduling of mixed-criticality systems must simultaneously deal with two contradictory goals:

- 1) On one hand, and since each task is supposed to carry out some useful computation for the system, it is desirable to respect the timeliness of all tasks, whatever their criticality is, which means that the over-provisioning of

resources required during the certification of the more critical tasks could be untighten somehow at run-time;

- 2) On the other hand, and since the higher the criticality of a task is, the more dramatic the consequences related to the failure of respecting their timeliness could be, it seems rather cautious to strictly respect the over-provisioning of resources of the more critical tasks, even if this means to interfere with the timeliness of the less critical ones (which is as meaningless as the criticality of these is low).

In practice, these goals are not necessarily exclusive though, and it is possible to pursue them both, at least to some extent. Indeed, beyond vouching for the correctness of the whole (or part of) the system, the certification process also allows for specifying for how long the functionalities can execute concurrently without interfering on each others temporal constraints. Indeed, when a task is certified by some CA, it is assuming a behavior, for the other tasks of the system, the pessimism of which depends on the criticality of the task(s) being certified. Consequently, as long as these assumptions hold at run-time, all the tasks that were certified up to that assurance level are guaranteed to meet their deadline. However, if a more critical task violates these assumptions, by executing longer than was assumed during the certification process, then the timeliness of the whole system can no longer be guaranteed anymore, which results in suspending the tasks that are less critical to ensure the timeliness of the tasks being more critical. In the state-of-the-art literature, the period during which these tasks are prevented from executing, henceforth called the *suspension delay*, is considered to be unlimited.

With little thought, one can notice that mixed-criticality systems can be seen as systems that exhibit *multiple behaviors*, issued from several operating modes, where each of these operating modes is characterized by a given set of functionalities. Such systems are commonly referred to, in the real-time literature, as *multi-moded* systems [1]. When viewing mixed-criticality systems as multi-moded systems, the transition from one operating mode to the other reduces the set of tasks performed by the system. Indeed, less critical tasks are suspended, but new tasks are never introduced (at least in the traditional way of handling mixed-criticality systems, as we will present an approach which allows to overstep this restriction). The transition from one mode to the other is triggered when the system detects that the certifications assumptions does not hold anymore.

When considering the temporal robustness of the system,

mixed-criticality is a natural approach, since it aims at favoring the timeliness of the functionalities that are crucial for the system. Nevertheless, discarding carelessly tasks that are less critical may have severe consequences on the system's consistency. Indeed, a task that is executing could be accessing data, the integrity and consistency of which has to be preserved. Under these assumptions, it is strongly undesirable to kill a job incautiously. It would therefore be wiser to allow the jobs released by less critical tasks to complete their execution when the transition from one operating mode to the other is triggered, but this transition stage might introduce a transient overload which should have no impact on the timeliness of high criticality tasks. Furthermore, while task suspension is carried out with concern for the respect of the timeliness of the high critical tasks, a suspended functionality can no longer handle the task it is in charge of. Consequently, an everlasting suspension delay might appear as strongly undesirable, and we might wish to be able to re-enable a suspended task at some point during the execution of the system.

A. Related Work

Mixed-Criticality scheduling is a research domain initially introduced by Vestal [2]. Nowadays, the Mixed-Criticality (MC)-Schedulability problem is known to arise in the context of applications being subject to multiple certification requirements. Many work has addressed the problem for systems implemented upon *uniprocessor* platforms [3]–[8]. More recently, research has been oriented towards the study of mixed-criticality scheduling upon multiprocessor platforms [9] [10]. Let us highlight that many of these work focused on a restricted and easier case of mixed-criticality systems, called dual-criticality systems, presenting only two distinct criticality levels. Nevertheless, to our knowledge, no work has tackled the problem that arises when mixed-criticality real-time tasks are incautiously discarded. Though in a previous work [8], we highlighted the fact that task suspension was often carried out too early, and that the system's computational resources could be exploited more cleverly. We also addressed the problem of managing task re-enablement. However, our work focused only on uniprocessor platforms, while this work will be applied to the more general case of identical multiprocessor platforms.

B. This Research

We believe that a multi-mode approach can bring some insight on how the problems mentioned in Section I-A could be formalized. To our knowledge, mixed-criticality systems have never formally been defined as multi-moded systems. In this research, we thus seek to establish a formal link between mixed-criticality and multi-moded hard real-time systems upon identical multiprocessor platforms. We will then show how to safely complete the execution of jobs released by the less critical tasks, potentially after their deadline, when the computational demand of the more critical tasks in the systems increases. In our opinion, it is safer in terms of system integrity to complete a job beyond its deadline, instead of dropping it carelessly once the latter is reached. We will also show

how to safely reduce the suspension delay suffered by the less critical tasks as the system load decreases, by allowing the re-enablement of the latter. This problem was almost never tackled in the current literature, but still presents a significant interest, since the less critical tasks can improve the overall behavior of the system. Finally, since the notion of mixed-criticality systems initially did not put any restriction on the number of criticality levels, and has been implemented using various degrees of criticality by industrial standards¹, our approach can be applied to task sets having any number of criticality level. This is an attractive feature, as it is much more compliant with current industrial standards.

II. MODEL AND DEFINITIONS

A. Mixed-criticality Specification

We consider multiprocessor platforms composed of a fixed number m of identical processors, denoted by $\pi = \{\pi_1, \pi_2, \dots, \pi_m\}$. Furthermore, we consider mixed-criticality sporadic task sets $\tau = \{\tau_1, \tau_2, \dots, \tau_n\}$, where the maximum criticality of a task is bounded by a natural value which we denote by L . A task τ_i in such a system is characterized by a 4-tuple of parameters $\{T_i, D_i, L_i, C_i\}$ where:

- $T_i \in \mathbb{N}_0$ is the minimum inter-arrival time separating two consecutive activations of task τ_i ;
- $D_i \in \mathbb{N}_0$ is the deadline of task τ_i , with $D_i \leq T_i$;
- $L_i \in \mathbb{N}_0$ is the criticality of the task τ_i , with $L_i \leq L$;
- $C_i \in \mathbb{N}_0^L$ is a size L vector of WCET, where $C_i(\ell)$ is an estimation of the WCET of task τ_i at criticality level $\ell \in [1, L]$.

Given these parameters, each task τ_i will generate a potentially infinite sequence of jobs, each release being separated by at least T_i time units, and each job having a hard deadline D_i time units after its release. We assume $C_i(\ell)$ is monotonically increasing for increasing values of ℓ . More precisely, for task τ_i the following two conditions hold:

- $\forall \ell \in [1, L_i), C_i(\ell) \leq C_i(\ell + 1)$;
- $\forall \ell \in [L_i, L], C_i(\ell) = C_i(L_i)$.

It follows that no task is supposed to execute longer than its WCET at its own criticality level. The k^{th} job $J_{i,k}$ released by a mixed-criticality task τ_i is characterized by a 3-tuple $\{r_{i,k}, d_{i,k}, c_{i,k}\}$ where:

- $r_{i,k} \in \mathbb{N}$ is the time instant at which $J_{i,k}$ was released. Since we consider sporadic task systems, we have $r_{i,k} - r_{i,k-1} \geq T_i$, and $r_{i,1} \geq 0$;
- $d_{i,k} \in \mathbb{N}_0$ is the absolute deadline of $J_{i,k}$. More precisely, $d_{i,k} \stackrel{\text{def}}{=} r_{i,k} + D_i$;
- $c_{i,k} \in \mathbb{N}_0$ is the exact execution time of $J_{i,k}$. From the specifications of τ_i , we can say that $c_{i,k} \leq C_i(L_i)$, but the exact value of $c_{i,k}$ will not be known until $J_{i,k}$ completes its execution;
- $f_{i,k} \in \mathbb{N}_0$ is the absolute time at which $J_{i,k}$ finishes its execution. Again, this value will only be known once $J_{i,k}$ actually completes its execution.

¹The RTCA-DO178B standard, for aerospace, defines 5 criticality levels, while the ISO26262, in the automotive industry, defines 4 criticality levels.

Definition 1 (Available job). At any time t , we call the k^{th} job $J_{i,k}$ released by task τ_i *available* if $t \geq r_{i,k}$ and $J_{i,k}$ has not yet completed its execution.

The actual execution time of job $J_{i,k}$ is not known from the specification of τ_i , but will only be discovered when $J_{i,k}$ completes its execution. Besides, the behavior of task τ_i might change from one execution to the other, so the actual execution times of the sequence of jobs released by τ_i may vary, leading to introduces the notion of *task scenario*.

Definition 2 (Task scenario). We define the scenario s_i^t of task τ_i at time t as the set of exact execution times $\{c_{i,1}, \dots, c_{i,k}\}$ for each of the k (k being dependent of t) jobs released by τ_i that already completed their execution at time t .

Definition 3 (Task set scenario). We define the scenario of the mixed-criticality task set τ at time t as $s^t = \{s_1^t, \dots, s_n^t\}$.

Definition 4 (ℓ -interval). Given a task set scenario, an ℓ -interval is an interval $[t_a, t_b)$ such that at time t_a , the system switched at criticality level ℓ , and the criticality of the system remained at level ℓ until time t_b .

Definition 5 (Worst-case response time). The worst-case response time (WCRT) $R_i(\ell)$ of task τ_i at criticality level ℓ is the maximum duration to execute any job of τ_i , in any scenario of criticality ℓ .

B. Modeling Mixed-Criticality in Terms of Multi-Mode

In this section, we propose a formalization of sporadic real-time mixed-criticality tasks in terms of multi-mode tasks. Each such task can be seen as being composed of different execution modes, and can be represented by means of a tuple, as described below:

$$\tau_i = \{\tau_i^{M_1} = \{T_i^{M_1}, D_i^{M_1}, C_i^{M_1}\}, \dots, \tau_i^{M_L} = \{T_i^{M_L}, D_i^{M_L}, C_i^{M_L}\}\}$$

where M_1, \dots, M_L represent the different operating modes to which the task belongs, and $X_i^{M_y}$ represents the value of parameter X of task τ_i when τ_i is executed in the operating mode M_y . Consequently, a natural representation for the mixed-criticality task $\tau_i = \{T_i, D_i, L_i, \{C_i(1), \dots, C_i(L_i)\}\}$ is by means of the following tuple:

$$\tau_i = \{\tau_i^1 = \{T_i, D_i, C_i(1)\}, \dots, \tau_i^{L_i} = \{T_i, D_i, C_i(L_i)\}\}$$

The interpretation of the tuple is the following: when the system is executing in operating mode M_{ℓ_1} , task τ_i is executed according to the parameters given by $\tau_i^{\ell_1}$. The scheduler thus makes the assumption that τ_i will not generate a job the execution time of which will exceed $C_i(\ell_1)$. If τ_i does release a job whose execution exceeds $C_i(\ell_1)$, then a transition from one operating mode to another occurs. If the system switches from mode M_{ℓ_1} to mode M_{ℓ_2} , then τ_i is executed according to the parameters given by $\tau_i^{\ell_2}$. From that specific time onward, the scheduler consequently assumes that τ_i will not generate jobs whose executions will exceed $C_i(\ell_2)$. A task τ_i belongs to every operating mode up to its own criticality level. More precisely, $\tau_i \in M_\ell \Leftrightarrow \ell \leq L_i$. A task τ_i will be enabled in every operating mode to which it belongs, and will be

suspended in every other operating mode.

Definition 6 (Enabled/suspended task). At run-time, a task τ_i is said to be enabled if τ_i can generate new jobs that are dispatched by the scheduler. Otherwise, τ_i is said to be suspended.

In the following, we will consider that the system is running is mode M_ℓ at time t if and only if all tasks of criticality less than ℓ are disabled, no job of criticality less than ℓ is still available at time t , and every task of criticality greater than or equal to ℓ is enabled. Notice that at system start-up, the system is running in operating mode M_1 , which implies that every task is enabled. In the remainder of the paper, we will denote by $\tau^\ell = \{\tau_{\ell_1}, \dots, \tau_{\ell_n}\} \subseteq \tau$ the set of tasks belonging to operating mode M_ℓ , i.e. the set of tasks τ_j with $L_j \geq \ell$.

C. Presentation of the MSM Scheduler

In Section II-A, we introduced the concept of criticality level of the scenario. This concept allows to define what is considered as being a feasible schedule when considering mixed-criticality systems.

Definition 7 (Feasible schedule). A schedule for a scenario s^t is feasible if, during every ℓ -interval, $1 \leq \ell \leq L$, every job $J_{i,k} | L_i \geq \ell$ ($1 \leq i \leq n$) that completed its execution by time t , received execution time $c_{i,k}$ between $r_{i,k}$ and $d_{i,k}$.

This definition implies that mixed-criticality scheduling is only concerned with respecting temporal constraints of tasks the criticality of which is higher than or equal to the criticality level of the scenario.

Definition 8 (\mathcal{S} -Schedulable). Let \mathcal{S} be a scheduling policy, and τ a mixed-criticality task set. We say τ is \mathcal{S} -Schedulable if, for *any* scenario of τ , \mathcal{S} generates a feasible schedule.

Since an on-line scheduling policy discovers the exact execution time of the jobs when they complete their execution, the criticality level of the scenario is not known beforehand. In the state-of-the-art way of handling mixed-criticality systems, as soon as a job exceeds its WCET of level ℓ , the criticality of the scenario is raised to level $\ell + 1$, all available jobs of criticality lower than $\ell + 1$ are *dropped*, and future releases from tasks of which the criticality is lower than $\ell + 1$ are no longer taken into consideration. At this point, recall that the goal of this research is to overstep these restrictions. Finally, the following definition specifies what is considered as being an MC-Schedulable system.

Definition 9 (MC-Schedulable). A mixed-criticality task set τ is MC-schedulable if there exists a scheduling policy \mathcal{S} such that τ is \mathcal{S} -schedulable.

In the remainder of this paper, we consider the multi-processor *global preemptive work-conservative static-priority* scheduler MSM proposed by Pathan [9], with the following interpretation:

- at each time instant, a global scheduler dispatches the m highest priority jobs (if any) on the m processors of the platform;
- a preemptive scheduler reserves the right to interrupt a job J_i belonging to a task τ_i , that is executing on a given processor, to assign this processor to another available

job J_j belonging to a task τ_j , with $i \neq j$;

- a work-conservative scheduling strategy never keeps a processor idle when there are available jobs;
- a static-priority scheduler assigns priorities to the tasks of the system, each job being scheduled inheriting from the priority of the task that released it.

The algorithm MSM thus dispatches available jobs according to traditional global static-priority scheduling, but has two additional implementation features, namely *job execution monitoring*, in order to detect a transition to a higher criticality level, and *task suspension*, to drop the less critical tasks when such a transition is detected. To our knowledge, MSM is the only multiprocessor scheduler that can be applied to mixed-criticality task sets having any number of criticality levels. Furthermore, and as already motivated by Pathan [9], static-priority schedulers are generally preferred by industries, as their decisions are predictable, thus enforcing the reliability property of the real-time systems they schedule.

The priority ordering of a task set is constructed according to Audsley's Optimal Priority Assignment (OPA) approach [11], according to the response time analysis of the system. These notions are further discussed in Sections II-C1 and II-C2. In the remainder of the paper, we will denote by $\text{hp}(\tau_i)$ the set of tasks τ_j having a priority higher than τ_i , and by $\text{lp}(\tau_i)$ the set of tasks τ_j having a priority lower than τ_i .

1) The Worst-Case Response Time Computation:

This section briefly introduces the work proposed by Pathan [9]. An upper-bound on the WCRT $R_i(\ell)$ of task τ_i at criticality level $\ell \in [1, L_i]$ is computed by considering the execution of any job $J_{i,k}$ released by τ_i in a window starting at $J_{i,k}$'s release time $r_{i,k}$, and finishing Δ time units later, i.e. at time instant $r_{i,k} + \Delta$. The procedure aims at defining an upper-bound on the *workload*, the *interfering workload*, the *total interfering workload*, and the *interference* (as explained in the following definitions) suffered by task τ_i in any scenario of criticality level ℓ during that interval.

Definition 10 (Workload). The *workload* of a higher priority task τ_j within the window of size Δ is the cumulative length of time interval during which the jobs released by τ_j execute within the window.

A task $\tau_j \in \text{hp}(\tau_i)$ is considered as a *carry-in* task if τ_j released a job before the start of the window, and that job executes (partially or fully) within the window. Otherwise, τ_j is considered as a *non carry-in* task. Pathan [9] highlighted the fact that if a higher-priority task is a carry-in task, then its worst-case interference on the lower-priority task is higher than if it was non carry-in.

Definition 11 (Carry-in/Non Carry-In Interfering Workload). The *carry-in interfering workload* $\bar{I}_{j,i}^{\text{CI}}(\Delta, \ell)$ (resp. *non carry-in interfering workload* $\bar{I}_{j,i}^{\text{NC}}(\Delta, \ell)$) of τ_j on task τ_i is the cumulative length of time interval during which a job released by a carry-in task τ_j (resp. a non carry-in task τ_j) executes, and $J_{i,k}$ not dispatched on any processor.

The difference between the carry-in and non carry-in interfering workload of a task $\tau_j \in \text{hp}(\tau_i)$ will be denoted by $\bar{I}_{j,i}^{\text{DIFF}}(\Delta, \ell) \stackrel{\text{def}}{=} \bar{I}_{j,i}^{\text{CI}}(\Delta, \ell) - \bar{I}_{j,i}^{\text{NC}}(\Delta, \ell)$

Definition 12 (Total Interfering Workload). The *total interfering workload* $\bar{I}_i(\Delta, \ell)$ is the sum of interfering workload of all the higher priority tasks within the window.

The MSM algorithm computes the total interfering workload of task τ_i as follows:

$$\bar{I}_i(\Delta, \ell) = \sum_{\tau_j \in \text{hp}(\tau_i)} \bar{I}_{j,i}^{\text{NC}}(\Delta, \ell) + \sum_{\tau_j \in \text{hp}_{m-1}(\tau_i)} \bar{I}_{j,i}^{\text{DIFF}}(\Delta, \ell) \quad (1)$$

where hp_{m-1} is the set of at most $m - 1$ carry-in tasks belonging to $\text{hp}(\tau_i)$ that have the largest value of $\bar{I}_{j,i}^{\text{DIFF}}(\Delta, \ell)$.

The reason to consider at most $m - 1$ carry-in tasks comes from a discussion from Guan *et al.* [12], formalized in the following property.

Property 1. The total interfering workload is upper-bounded by considering at most $m - 1$ carry-in tasks within the window of any lower priority task, when considering global static priority scheduling of constrained-deadline sporadic task sets.

Definition 13 (Interference). The interference suffered by a task τ_i in any scenario of criticality level ℓ , and during a time interval of length Δ , is the cumulative length of time interval during which the m processors are busy executing tasks belonging to $\text{hp}(\tau_i)$. An upper-bound on the interference suffered by τ_i over an interval of length is given by $\left\lfloor \frac{\bar{I}_i(\Delta, \ell)}{m} \right\rfloor$.

Finally, and from the above definitions, since in any scenario of criticality level ℓ , $J_{i,k}$ is allowed to execute for at most $C_i(\ell)$ time units, the WCRT $R_i(\ell)$ of task τ_i at criticality level ℓ is obtained by determining the least fixed point of the following function:

$$R_i(\ell) = C_i(\ell) + \left\lfloor \frac{\bar{I}_i(R_i(\ell), \ell)}{m} \right\rfloor \quad (2)$$

Since $\bar{I}_i(\Delta, \ell)$ is an upper-bound on the total interfering workload suffered by task τ_i , the *actual* total interfering workload suffered by τ_i over an interval of length Δ , in any scenario of criticality level ℓ , will be denoted by $I_i^*(\Delta, \ell)$.

2) Finding Priorities Using Audsley's Approach:

The response time analysis described in Section II-C1 is used to find a static-priority ordering of the mixed-criticality task set τ , as depicted by Algorithm 1. Indeed, at each step, the method tries to identify a task τ_i satisfying $R_i(\ell) \leq D_i$, $1 \leq \ell \leq L_i$ ². If such a task is found, then the same reasoning is iteratively applied to task set $\tau \setminus \{\tau_i\}$. Otherwise, if for every task τ_i , $\exists \ell | R_i(\ell) > D_i$, then the method fails.

D. Mode Transition Specifications

A change in the operating mode of the system is instantiated whenever the system detects a change in its internal state. More precisely, switching from criticality level ℓ_1 to criticality level ℓ_2 can be seen as switching from one operating mode, referred to as the *old-mode* M_{ℓ_1} , to another operating mode, referred to as the *new-mode* M_{ℓ_2} . A *Mode Change Request* (MCR) is defined as being the event that triggers a mode change transition. In the current state-of-the-art literature, the only MCR that is considered is triggered whenever the system

²This is due to the fact that to our knowledge, there is no proof that $R_i(\ell) \leq R_i(\ell + 1)$, $1 \leq \ell < L$ holds.

Algorithm 1: OPA Algorithm

```
pr  $\leftarrow |\tau|$ 
while  $\tau \neq \emptyset$  do
  Let  $\tau_i$  be a task from  $\tau$ ;
  if  $R_i(\ell) \leq D_i$ ,  $1 \leq \ell \leq L_i$  then
    Assign  $\tau_i$  the priority pr;
     $\tau \leftarrow \tau \setminus \{\tau_i\}$ ;
    pr  $\leftarrow$  pr - 1;
  else
    return error;
```

detects that a task τ_i exceeds its WCET at criticality level ℓ , $1 \leq \ell < L_i$, resulting in the system to switch from operating mode M_ℓ to operating mode $M_{\ell+1}$. In the remainder of this paper, this MCR will be referred to as an *increasing mode change request* denoted by $\text{IMCR}_{\ell+1}$.

Definition 14 (Increasing Mode Change Request). Whenever the system is running in operating mode M_ℓ , an increasing mode change request from criticality level ℓ to criticality level $\ell + 1$, $\forall \ell \in [1, L]$, denoted by $\text{IMCR}_{\ell+1}$, is the event that will result in switching from operating mode M_ℓ to operating mode $M_{\ell+1}$. The time at which $\text{IMCR}_{\ell+1}$ is triggered will be denoted by $t_{\text{IMCR}_{\ell+1}}$.

To handle this transition, the traditional approach that is currently implemented in the state-of-the-art literature consists in suspending *instantaneously* and *forever* tasks that do not belong to the new mode, their potential available jobs at time $t_{\text{IMCR}_{\ell+1}}$, henceforth called the rem-jobs, not even being dispatched anymore. The transition stage is thus trivial, and the system instantaneously switches from operating mode M_ℓ to operating mode $M_{\ell+1}$. By suspending the less critical tasks upon an $\text{IMCR}_{\ell+1}$, the goal is to respect both the $(\ell+1)$ -periodicity and $(\ell+1)$ -feasibility, defined as follows.

Definition 15 (ℓ -periodicity, adapted from [1]). The ℓ -periodicity is the property that requires the activation pattern of each task τ_i such that $L_i \geq \ell$ to be respected. In other words, the activation pattern of a task τ_i should not be altered by increasing mode change requests up to τ_i 's own criticality.

Definition 16 (ℓ -feasibility, adapted from [1]). The ℓ -feasibility is the property that requires each task τ_i such that $L_i \geq \ell$ to meet its deadline.

Consequently, an increasing mode change request $\text{IMCR}_{\ell+1}$ $\forall \ell < L_i$ should have no impact on the deadline of a task τ_i . With concern for the respect of these properties, the traditional approach mentioned above however introduces two major drawbacks. On one hand, and since the tasks belonging to the old-mode might be accessing data, it might appear more cautious to allow them to complete their execution before discarding them. This nevertheless could result in a temporary overload which could jeopardize the timeliness of the new-mode tasks. The transition should therefore be managed in such a way that the $(\ell+1)$ -periodicity and $(\ell+1)$ -feasibility properties are preserved, while the rem-jobs could be allowed to complete their execution. Furthermore, as an everlasting suspension might not be necessary to keep on respecting the

$(\ell+1)$ -periodicity and $(\ell+1)$ -feasibility, we would also like to re-enable the less critical tasks as soon as possible. Achieving the second goal nevertheless calls for the use of a new mode change request, namely a *decreasing mode change request*.

Definition 17 (Decreasing Mode Change Request). Whenever the system is running in operating mode M_h , a decreasing mode change request from criticality level h to ℓ , $\forall \ell \in [1, h)$, denoted by DMCR_ℓ , is the event that will result in switching from operating mode M_h to operating mode M_ℓ . The time at which DMCR_ℓ is triggered will be denoted by t_{DMCR_ℓ} .

We again insist on the fact that current research in the field of mixed-criticality never considers the re-enablement of previously suspended tasks, at the exception of [8]. Based on the two mode change requests that are considered, we present in Section III mode change protocols whose goals are:

- Upon an $\text{IMCR}_{\ell+1}$, complete the execution of the rem-jobs as soon as possible, as this will improve system consistency;
- Upon a DMCR_ℓ , re-enable the previously suspended tasks belonging to mode M_ℓ as soon as possible, as this will improve the system's overall behavior.

Definition 18 (Valid increasing transition protocol). Assuming the system is running in operating mode M_ℓ , an increasing transition protocol \mathcal{P} is said to be *valid* if and only if, upon an $\text{IMCR}_{\ell+1}$, \mathcal{P} ensures both the $(\ell+1)$ -periodicity and $(\ell+1)$ -feasibility.

Definition 19 (Valid decreasing transition protocol). Assuming the system is running in operating mode M_h , a decreasing mode transition protocol \mathcal{P} is said to be *valid* if and only if, upon an DMCR_ℓ , \mathcal{P} ensures both the h -periodicity and h -feasibility.

It follows that the traditional approach is a valid increasing transition protocol, since upon a $\text{IMCR}_{\ell+1}$, and from time $t_{\text{IMCR}_{\ell+1}}$ onward, it aims only at respecting the timeliness of tasks the criticality of which is at least equal to $\ell + 1$. Furthermore, the former approach does not consider decreasing mode change requests, so the timeliness of the high critical tasks can not be jeopardized by the re-enablement of the less critical tasks. By extending this approach to allow the rem-jobs to complete their execution, and allow the re-enablement of the less critical tasks, we thus only seek to achieve a cleverer usage of the platform on which the mixed-criticality system is executed, without compromising the reliability of the considered mixed-criticality system. In the literature related to multi-moded systems, it is sometimes assumed that a mode change request $\text{IMCR}_{\ell+1}$ may only be triggered in the steady state of the system, i.e. whenever the system is in mode M_ℓ . In the case of mixed-criticality systems however, a mode change request $\text{IMCR}_{\ell+2}$ could happen during the handling of the mode change request $\text{IMCR}_{\ell+1}$. Indeed, in mode M_ℓ , it is assumed no task τ_i will generate a job the execution time of which will exceed $C_i(\ell)$. If this nevertheless happens, then $\text{IMCR}_{\ell+1}$ is triggered, but this does not prevent the job of τ_i to keep on executing, potentially leading to exceed $C_i(\ell+1)$ and trigger $\text{IMCR}_{\ell+2}$ while still handling $\text{IMCR}_{\ell+1}$.

III. THE PROTOCOL

A. Introductory Definitions

The OPA procedure described in Section II-C considers specific scenarios for each tasks of the system, where the execution time of each job $J_{i,k}$ is exactly equal to the WCET of task τ_i at a given criticality level ℓ such that $1 \leq \ell \leq L_i$. These scenarios are called *basic*, as formalized by the following definitions.

Definition 20 (Basic task scenario). At any time t , we call the scenario s_i^t of task τ_i basic if for all $c_{i,k} \in s_i^t$, there exists a criticality level $\ell_k \in [1, L_i]$ such that $c_{i,k} = C_i(\ell_k)$.

In a basic task scenario, each job released by a task τ_i will thus complete its execution after having consumed exactly its WCET at a given criticality level (which is not necessarily equal to the task's own criticality level L_i).

Definition 21 (Basic task set scenario). At any time t , we call the scenario of the mixed-criticality task set τ basic if every $s_i^t \in s^t$ is a basic task scenario.

Definition 22 (Basically \mathcal{S} -Schedulable). Let \mathcal{S} be a scheduling policy, and τ a mixed-criticality task set. We say τ is basically \mathcal{S} -schedulable if, for *any* basic scenario of τ , \mathcal{S} generates a feasible schedule.

Theorem 1 (Baruah *et al.* [13]). A mixed-criticality task set τ is MC-schedulable on a given platform π if τ is basically \mathcal{S} -schedulable by a given scheduling policy \mathcal{S} .

The OPA procedure thus aims at determining whether a given mixed-criticality task set τ is basically MSM-schedulable. Indeed, during the WCRT computation at criticality level ℓ , $\forall \ell \in [1, L]$, the procedure considers each task τ_i will release jobs whose execution times will be equal to τ_i 's WCET a some criticality level. The scenarios that are considered are thus basic. Consequently, if the OPA procedure succeeds, τ is deemed basically MSM-schedulable, and according to Theorem 1, τ is also MC-schedulable.

Throughout this section, we will consider that the priority of a task is given by its index, such that task τ_i has a higher priority than task τ_{i+1} .

B. Handling Increasing Mode Change Requests

The offline analysis stage, whose goal is to vouch for the feasibility of the considered mixed-criticality system, introduces two sources of pessimism:

- The WCET estimation introduces a level of pessimism which is as significant as the criticality of the considered task is high;
- The WCRT computation, as it is carried out in Section II-C1, leads to defining an *upper-bound* on the WCRT of each task, rather than the actual *exact* WCRT.

The introduction of these sources of pessimism leads to overestimate the actual requirements of the system. In particular, this means that during the execution of the system, a lot of computation resources could be wasted. In this section, we will show how to take advantage of the pessimism introduced during both the WCET and WCRT analysis stages by reclaiming these wasted resources, to allow the less critical tasks to

complete their execution during the online scheduling phase. Before diving in the details of our contributions, recall that the WCRT computation assumes that whenever the system switches from operating mode M_ℓ to operating mode $M_{\ell+1}$ at time $t_{\text{IMCR}_{\ell+1}}$, each task $\tau_i \in \tau$ such that $L_i \leq \ell$, do not interfere anymore on tasks $\tau_j \in \tau$, such that $L_j > \ell$, from time $t_{\text{IMCR}_{\ell+1}}$ onward. To prevent this interference when switching from operating mode M_ℓ to operating mode $M_{\ell+1}$, the mixed-criticality system model prescribes the immediate suspension of each task satisfying $L_i \leq \ell$, their potential rem-jobs being instantaneously dropped as soon as the mode change request $\text{IMCR}_{\ell+1}$ is triggered. As a result, if the main goal is to avoid jeopardizing the timeliness of the tasks the criticality of which is at least equal to $\ell + 1$, the completion of the rem-jobs must be handled in such a way that no interference can occur on a task with a criticality at least equal to $\ell + 1$. We now present our contributions, by suggesting three alternative approaches to avoid this interference.

1) A Naive Approach:

Upon an $\text{IMCR}_{\ell+1}$, an intuitive approach consists in assigning the rem-jobs the criticality of which is less than $\ell + 1$ a priority that is *lower* than the jobs the criticality of which is at least equal to $\ell + 1$. The execution of the rem-jobs thus takes place whenever there are less than m available jobs released by tasks $\tau_j \in \tau^{\ell+1}$. The following theorem proves this approach to be correct, i.e. $(\ell + 1)$ -periodicity and $(\ell + 1)$ -feasibility is preserved among an $\text{IMCR}_{\ell+1}$.

Theorem 2. Upon an $\text{IMCR}_{\ell+1}$, assigning the rem-jobs of criticality level ℓ a priority that is lower than every task τ_i such that $L_i \geq \ell + 1$ is a valid increasing transition protocol.

Proof: By assigning the rem-jobs a priority that is lower than every task τ_i such that $L_i \geq \ell + 1$, they cannot interfere on the tasks $\tau_j \in \tau^{\ell+1}$. It follows that the $(\ell + 1)$ -periodicity and $(\ell + 1)$ -feasibility are preserved, and that this increasing transition protocol is valid. ■

2) A WCET-Based Approach:

The drawback of the naive approach lies in the fact that upon an $\text{IMCR}_{\ell+1}$, the completion of the rem-jobs having a criticality equal to ℓ can only occur whenever there are strictly less than m available jobs of criticality at least equal to $\ell + 1$. As a consequence, the rem-jobs could complete their execution well beyond their deadline. We can however take advantage of the pessimism introduced during the WCET analysis to allow the rem-jobs to complete their execution earlier. Indeed, recall that the offline analysis vouched for the feasibility of any basic scenario of the system. More precisely, and considering that the system is currently executing in operating mode M_ℓ , this means that the scheduler will make the assumption that every job $J_{i,k}$ will complete its execution after exactly $c_{i,k} = C_i(\ell)$ time units. Nevertheless, the fact is that $J_{i,k}$ could complete its execution after an amount of time $c_{i,k}$ such that $c_{i,k} \leq C_i(\ell)$, in any scenario of criticality ℓ . Let us therefore denote by $\text{ur}_{i,k}(\ell) \stackrel{\text{def}}{=} C_i(\ell) - c_{i,k}$ the unused resources of $J_{i,k}$ in any scenario of criticality level ℓ . The value $\text{ur}_{i,k}(\ell)$ thus represents a fraction of the resources that were reserved by the scheduler for job $J_{i,k}$, under the assumption of a basic

scenario of criticality ℓ , and that $J_{i,k}$ did not consume. Following from the above discussion, we propose an alternative increasing transition protocol, that slightly adapts the MSM scheduler during the transition phase. Indeed, the latter would act as if $J_{i,k}$ did not complete its execution after $c_{i,k}$ time units, and take advantage of the unused resources to carry on the execution of the rem-jobs during exactly $\text{ur}_{i,k}(\ell)$ time units. Upon an $\text{IMCR}_{\ell+1}$, recovering these unused resources to carry on the execution of the rem-jobs of criticality equal to ℓ will preserve both the $(\ell+1)$ -periodicity and $(\ell+1)$ -feasibility, as proved by the following theorem.

Theorem 3. Upon an $\text{IMCR}_{\ell+1}$, the protocol that reclaims the unused resources $\text{ur}_{i,k}(\ell+1)$ of every job $J_{i,k}$ of criticality at least equal to $\ell+1$, to complete the execution of the rem-jobs of criticality ℓ , is a valid increasing transition protocol.

Proof: The proof relies on the observation that reclaiming exactly $\text{ur}_{i,k}(\ell+1)$ units of computation from task τ_i consists in simulating a basic task scenario for τ_i . Indeed, $c_{i,k} + \text{ur}_{i,k}(\ell+1) = C_i(\ell+1)$. Since the system was deemed MSM-schedulable, it follows that any basic scenario is feasible. As a consequence, the proposed increasing transition protocol is valid. ■

3) A WCRT-Based Approach:

As explained in Section II-C1, an upper-bound on the WCRT $R_i(\ell)$ of task $\tau_i \in \tau$ at criticality level ℓ is computed assuming an upper-bound on the interference suffered by τ_i in the most pessimistic scenario of criticality level ℓ . Furthermore, in the most pessimistic scenario of criticality level ℓ , the job $J_{i,k}$ should execute for exactly $C_i(\ell)$ which is also an upper-bound on the execution time of task τ_i . Due to this twofold source of pessimism, it is consequently most unlikely that a job $J_{i,k}$ released by task τ_i will ever complete its execution exactly $R_i(\ell)$ time units after $r_{i,k}$. However, the system having been deemed MSM-schedulable, it follows that the timeliness of each of its tasks could be respected in presence of this pessimism. In this section, we will thus suggest an increasing transition protocol which will enable for the reclaiming of unused resources, this time being based on the WCRT of each task. Our approach is based on the fact that the offline analysis assumed that a job $J_{i,k}$ could complete its execution $R_i(\ell)$ time units after $r_{i,k}$ in any scenario of criticality level ℓ . If, however, $J_{i,k}$ completes its execution earlier, the scheduler will simulate the availability of $J_{i,k}$ until time $r_{i,k} + R_i(\ell)$, by assigning a processor to the rem-jobs whenever $J_{i,k}$ would have been executed. The challenge of this section is to prove that this will not increase the WCRT of tasks having a lower priority. The following theorem formally proves that if every job released by tasks $\tau_j \in \text{hp}(\tau_i)$ completes its execution no later than $R_j(\ell)$ time units after its release, then every job released by task τ_i will complete its execution no later than $R_i(\ell)$ time units after its release as well, in any scenario of criticality level ℓ .

Theorem 4. Each job $J_{i,k}$ released by a task τ_i will complete its execution no later than $R_i(\ell)$ time units after its release in any scenario of criticality level $\ell \leq L_i$, if for every task $\tau_j \in \text{hp}(\tau_i)$, every job released by a task τ_j completes its

execution no later than $R_j(\ell)$ time units after its release.

Proof: The proof is by induction on the priorities of the tasks. We will indeed show that if the property holds for every task up to τ_{i-1} , then it must also hold for task τ_i .

Base case: The base case consists in considering task τ_1 , i.e. the task having the highest priority. Since τ_1 suffers no interference at all, it is trivial that every job released by τ_1 will complete its execution no later than $R_1(\ell) = C_1(\ell)$ time units after its release, in any scenario of criticality level ℓ .

Induction step: Assume that the property holds for every task up to τ_{i-1} and let us show that it must also hold for τ_i , by considering the execution of any job $J_{i,k}$ released by τ_i over an interval of length Δ . We will distinguish between two cases:

Case 1: let us first consider that there are no more than $m-2$ carry-in tasks belonging to $\text{hp}(i-1)$ over the interval of length Δ . From Equation 1, according to the inductive hypothesis and since there are at most $m-2$ carry-in tasks belonging to $\text{hp}(\tau_{i-1})$, the actual total interfering workload suffered by task τ_{i-1} over the interval of length Δ can be expressed as follows:

$$\begin{aligned} \mathbb{I}_{i-1}^*(\Delta, \ell) &\leq \sum_{\tau_j \in \text{hp}(\tau_{i-1})} \bar{\mathbb{I}}_{j,i-1}^{\text{NC}}(\Delta, \ell) \\ &\quad + \sum_{\tau_j \in \text{hp}_{m-2}(\tau_{i-1})} \bar{\mathbb{I}}_{j,i-1}^{\text{DIFF}}(\Delta, \ell) \end{aligned} \quad (3)$$

Because carry-in tasks have a higher interfering workload than non-carry-in tasks, an upper-bound on the interfering workload cause by task τ_{i-1} on task τ_i over the interval of length Δ is given by the following equation:

$$\mathbb{I}_{i-1,i}^*(\Delta, \ell) \leq \bar{\mathbb{I}}_{i-1,i}^{\text{NC}}(\Delta, \ell) + \bar{\mathbb{I}}_{i-1,i}^{\text{DIFF}}(\Delta, \ell) \quad (4)$$

Since the total interfering workload suffered by τ_i is equal to the total interfering workload suffered by τ_{i-1} , plus the interfering workload suffered by τ_i from τ_{i-1} , using Equations 3 and 4, the actual total interfering workload suffered by job $J_{i,k}$ over the interval of length Δ can be upper-bounded as follows:

$$\mathbb{I}_i^*(\Delta, \ell) \leq \mathbb{I}_{i-1}^*(\Delta, \ell) + \bar{\mathbb{I}}_{i,i-1}^{\text{NC}}(\Delta, \ell) + \bar{\mathbb{I}}_{i,i-1}^{\text{DIFF}}(\Delta, \ell) \quad (5)$$

However:

$$\begin{aligned} \bar{\mathbb{I}}_{i,i-1}^{\text{DIFF}}(\Delta, \ell) &+ \sum_{\tau_j \in \text{hp}_{m-2}(\tau_{i-1})} \bar{\mathbb{I}}_{j,i-1}^{\text{DIFF}}(\Delta, \ell) \\ &\leq \sum_{\tau_j \in \text{hp}_{m-1}(\tau_i)} \bar{\mathbb{I}}_{j,i}^{\text{DIFF}}(\Delta, \ell) \end{aligned} \quad (6)$$

Using Equation 1, we are able to conclude that:

$$\begin{aligned} \mathbb{I}_i^*(\Delta, \ell) &\leq \sum_{\tau_j \in \text{hp}(\tau_i)} \bar{\mathbb{I}}_{j,i}^{\text{NC}}(\Delta, \ell) \\ &\quad + \sum_{\tau_j \in \text{hp}_{m-1}(\tau_i)} \bar{\mathbb{I}}_{j,i}^{\text{DIFF}}(\Delta, \ell) = \bar{\mathbb{I}}_i(\Delta, \ell) \end{aligned} \quad (7)$$

Equation 7 proves that the actual total interfering workload suffered by $J_{i,k}$ will not exceed $\bar{\mathbb{I}}_i(\Delta, \ell)$, thus implying that $J_{i,k}$ will complete its execution no later than $R_i(\ell)$ time units after $r_{i,k}$.

Case 2: let us now consider the case where there are *at least* $m-1$ carry-in tasks belonging to $\text{hp}(i-1)$ over the interval of length Δ . From Property 1, since the total interfering workload is upper-bounded when considering at most $m-1$ carry-in tasks, τ_{i-1} can be considered as a non carry-in task. An upper-bound on the interfering workload suffered by $J_{i,k}$ from task τ_{i-1} is thus given by:

$$\bar{I}_{i-1,i}(\Delta, \ell) \leq \bar{I}_{i-1,i}^{\text{NC}}(\Delta, \ell)$$

As a consequence, the actual total interfering workload suffered by job $J_{i,k}$ over the interval of length Δ , can be expressed as:

$$I_i^*(\Delta, \ell) \leq I_{i-1}^*(\Delta, \ell) + \bar{I}_{i-1,i}^{\text{NC}}(\Delta, \ell) \quad (8)$$

According to the inductive hypothesis, we have:

$$I_i^*(\Delta, \ell) \leq \sum_{\tau_j \in \text{hp}(\tau_i)} \bar{I}_{j,i}^{\text{NC}}(\Delta, \ell) + \sum_{\tau_j \in \text{hp}_{m-1}(\tau_{i-1})} \bar{I}_{j,i}^{\text{DIFF}}(\Delta, \ell) \quad (9)$$

However, since $\text{hp}(\tau_{i-1}) \subset \text{hp}(\tau_i)$, we have:

$$\sum_{\tau_j \in \text{hp}_{m-1}(\tau_{i-1})} \bar{I}_{j,i}^{\text{DIFF}}(\Delta, \ell) \leq \sum_{\tau_j \in \text{hp}_{m-1}(\tau_i)} \bar{I}_{j,i}^{\text{DIFF}}(\Delta, \ell) \quad (10)$$

By injecting Equation 10 into Equation 9, and by using Equation 1 we get the following upper-bound on the actual total interfering workload suffered by $J_{i,k}$:

$$I_i^*(\Delta, \ell) \leq \sum_{\tau_j \in \text{hp}(\tau_i)} \bar{I}_{j,i}^{\text{NC}}(\Delta, \ell) + \sum_{\tau_j \in \text{hp}_{m-1}(\tau_i)} \bar{I}_{j,i}^{\text{DIFF}}(\ell) = \bar{I}_i(\Delta, \ell) \quad (11)$$

Equation III-B3 proves that the actual total interfering workload suffered by $J_{i,k}$ will not exceed $\bar{I}_i(\Delta, \ell)$, implying that $J_{i,k}$ will complete its execution after no more than $R_i(\ell)$ time units. ■

Corollary 1. Each job $J_{i,k}$ released by a task τ_i will complete its execution no later than $R_i(\ell)$ time units after $r_{i,k}$ in any scenario of criticality level $\ell \leq L_i$, if every job $J_{j,p}$ released by a task $\tau_j \in \text{hp}(\tau_i)$ completes its execution exactly $R_j(\ell)$ time units after $r_{j,p}$.

From Corollary 1, we propose an alternative increasing transition protocol, that slightly adapts the MSM scheduler during the transition phase resulting from the triggering of $\text{IMCR}_{\ell+1}$. Upon the completion of a job $J_{i,k}$ released by a task τ_i such that $L_i \geq \ell+1$ at time $f_{i,k} < r_{i,k} + R_i(\ell+1)$, the scheduler would simulate the availability of $J_{i,k}$ in the interval $[f_{i,k}, r_{i,k} + R_i(\ell+1)]$, thus acting as if $J_{i,k}$ had suffered an interference equal to the upper-bound $I_i^\ell(R_i(\ell))$, and had completed its execution exactly at time $r_{i,k} + R_i(\ell)$. The following theorem proves that this approach will preserve both the $(\ell+1)$ -periodicity and $(\ell+1)$ -feasibility.

Theorem 5. Upon an $\text{IMCR}_{\ell+1}$, the protocol that simulates the availability of each job $J_{i,k}$ such that $L_i \geq \ell+1$ until time $r_{i,k} + R_i(\ell+1)$, to complete the execution of the rem-jobs of

criticality ℓ , is a valid increasing transition protocol.

Proof: This is an immediate consequence of Theorem 4. Indeed, by simulating the availability of each job $J_{i,k}$ released by task τ_i until time $r_{i,k} + R_i(\ell+1)$, we simulate the upper-bound on the interference suffered by $J_{i,k}$. According to Corollary 1, this will not increase the WCRT of tasks $\tau_j \in \text{lp}(\tau_i)$. As a consequence, the proposed increasing transition protocol is valid. ■

4) Conclusion and Discussion:

In this section, we have proposed several approaches to handle the transition from criticality level ℓ to criticality level $\ell+1$. The proposed approaches have a common goal, as they all allow to reach a compromise between a high safety, which is implied by the high level of pessimism adopted during the offline analysis phase, and an efficient usage of the platform, by recovering the unused resources to execute the less critical tasks during the online phase. In our opinion, this is an important aspect of our work, as it highlights the crucial fact that a high level of assurance does not necessarily imply that resources are doomed to be wasted.

Finally, note that the proposed approaches do not assume anything on the execution order of the rem-jobs. We will thus briefly discuss what management policy can be implemented in the case the system switches from operating mode M_ℓ to operating mode $M_{\ell+1}$ at time $t_{\text{IMCR}_{\ell+1}}$ while some rem-jobs of criticality equal to $\ell-1$ are still available. This means that at time $t_{\text{IMCR}_{\ell+1}}$, the system has to complete the execution of rem-jobs the criticality of which is less than ℓ . In that case, either it is assumed that since their criticality is less than the current criticality of the system, they are all equal in terms of importance, meaning that a global strategy can be applied to complete their execution (complete the rem-jobs with the shortest deadline first, complete the rem-jobs with the shortest remaining processing time first, etc.). Or we can assume that even though the current criticality of the system is higher than their own criticality, the relative importance of the rem-jobs remains the same, meaning it is safer to complete the rem-jobs in decreasing order of criticality level (the rem-jobs with a criticality equal to ℓ have to be completed before the rem-jobs with a criticality equal to $\ell-1$).

C. Handling Decreasing Mode Change Requests

In our previous work [8], we focused on *uniprocessor* platforms, and proved that whenever an idle time was detected, while the system had reached criticality level $\ell > 1$, it was safe to re-enable *every* task that had previously been suspended. However, the occurrence of a simultaneous idle time on every processors of a multiprocessor platform is unlikely. In this section, we will therefore suggest an alternative approach to re-enable the suspended tasks without relying on the occurrence of idle times. Before going any further, let us introduce the following property, which highlights two important requirements regarding task re-enablement.

Property 2. When the system is switching from operating

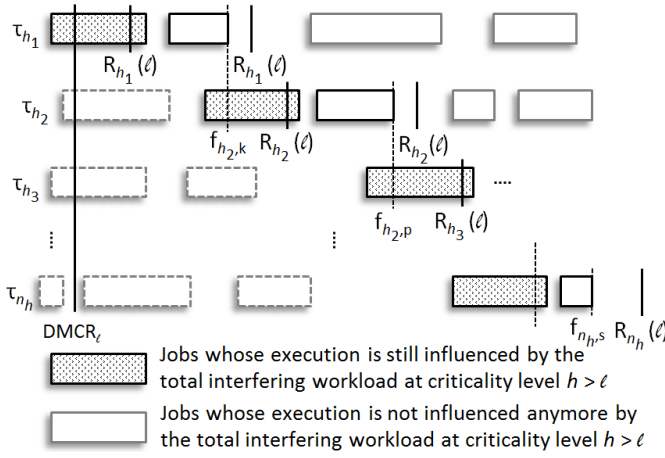


Figure 1. The decreasing mode transition protocol iteratively identifies the time instants $f_{i,k} \leq R_i(\ell)$ for every task $\tau_i \in \tau^h$. mode M_h to operating mode M_ℓ upon a DMCR_ℓ at time t_{DMCR_ℓ} , the re-enablement of a previously suspended task τ_i belonging to M_ℓ at time $t \geq t_{\text{DMCR}_\ell}$ may be carried out provided that the following two conditions hold:

- The re-enablement at time t of task τ_i must not jeopardize the h -periodicity and h -feasibility of the system;
- Task τ_i must be guaranteed to meet its deadline from time t onward if no $\text{IMCR}_{\ell+1}$ occurs at a time $t_{\text{IMCR}_{\ell+1}} \geq t$.

In the rest of this paper, we will say that it is *safe* to re-enable a task τ_i if Property 2 holds upon the re-enablement of τ_i . The way new-mode tasks are re-enabled leads to distinguish two types of protocols.

Definition 23 (Synchronous/Asynchronous protocol [1]). Assuming the system is switching from operating mode M_h to operating mode M_ℓ , a *synchronous* protocol is a protocol that re-enables each suspended task belonging to operating mode M_ℓ simultaneously. An *asynchronous* protocol is a protocol that enables every suspended task belonging to operating mode M_ℓ independently from the others, i.e. some tasks belonging to M_ℓ might be enabled earlier than others.

Assuming the system is executing in operating mode M_h , and depending on the set of functionalities the system is willing to re-enable, once the rem-jobs of criticality less than h are completed, a DMCR_ℓ , $1 \leq \ell < h$ can be triggered. For the sake of clarity, we will assume that no $\text{IMCR}_{\ell+1}$ is triggered during the decreasing mode change (we will shortly discuss this potential event later). We suggest a synchronous protocol, illustrated by Figure 1, that works as follows: starting at time t_{DMCR_ℓ} , the protocol identifies the first job $J_{h1,k}$ released by task τ_{h1} that completes its execution at time $f_{h1,k} \leq r_{h1,k} + R_{h1}(\ell)$. From time $f_{h1,k}$, the protocol then identifies the first job $J_{h2,p}$ released by task τ_{h2} that completes its execution at time $f_{h2,p} \leq r_{h2,p} + R_{h2}(\ell)$. The procedure then keeps on identifying one such job for each task $\tau_i \in \tau^h$, in order of their priority (i.e. it identifies such a job for task τ_{hi} only when it has previously identified such a job for task τ_{hi-1}). The re-enablement of the previously suspended tasks having a criticality at least equal to ℓ can then take place when the procedure identifies a job $J_{nh,s}$ that completes its

execution at time $f_{nh,s} \leq r_{nh,s} + R_{nh}(\ell)$. In the following, we will assume that if the protocol identifies a job $J_{hi,k}$ that completes its execution at time $f_{hi,k} \leq r_{hi,k} + R_i(\ell)$, then this means that the procedure has already identified such a job for every tasks $\tau_{h1}, \tau_{h2}, \dots, \tau_{hi-1}$.

Lemma 1. Whenever a task $\tau_{hi} \in \tau^h$ releases a job $J_{hi,k}$ that completes its execution at time $f_{hi,k} \leq r_{hi,k} + R_{hi}(\ell)$, and no $\text{IMCR}_{\ell+1}$ is triggered during the mode change, then the actual interfering workload suffered by tasks $\tau_j \in \text{lp}(\tau_{hi})$ from task τ_{hi} , from time $f_{hi,k}$ onward will be less than or equal to $\bar{I}_{hi,j}^{\text{CI}}(\Delta, \ell)$ (resp. $\bar{I}_{hi,j}^{\text{NC}}(\Delta, \ell)$) over any window of length Δ , if τ_{hi} is a carry-in (resp. non carry-in) task for τ_j .

Proof: The upper-bound on the interfering workload suffered by task $\tau_j \in \text{lp}(\tau_{hi})$ from task τ_{hi} over an interval of length Δ , in any ℓ -interval, is computed assuming τ_{hi} will release jobs that will execute for no more than $C_{hi}(\ell)$ time units. However, if no $\text{IMCR}_{\ell+1}$ is triggered, then it must be the case that every job released by task τ_{hi} executes for no more than $C_{hi}(\ell)$ time units. Therefore, from $f_{hi,k}$ onward, the actual total interfering workload suffered by tasks $\tau_j \in \text{lp}(\tau_{hi})$ from task τ_{hi} will be less than or equal to $\bar{I}_{hi,j}^{\text{CI}}(\Delta, \ell)$ (resp. $\bar{I}_{hi,j}^{\text{NC}}(\Delta, \ell)$) over any window of length Δ if τ_{hi} is a carry-in (resp. non carry-in) task for τ_j . ■

Lemma 2. Let us assume the protocol identified a job $J_{hi-1,k}$ that completed its execution at time $f_{hi-1,k} \leq r_{hi-1,k} + R_{hi-1}(\ell)$. From time $f_{hi-1,k}$ onward the actual total interfering workload $\bar{I}_i^*(\Delta, \ell)$ suffered by task τ_{hi} over any window of size Δ will be less than or equal to $\bar{I}_i(\Delta, \ell)$.

Proof: From Lemma 1, we know that from time $f_{hi-1,k}$ onward, the actual interfering workload of every task $\tau_j \in \text{hp}(\tau_{hi})$ on τ_i is less than or equal to $\bar{I}_{j,hi}^{\text{CI}}(\Delta, \ell)$ or $\bar{I}_{j,hi}^{\text{NC}}(\Delta, \ell)$, depending whether τ_j is a carry-in task for τ_{hi} or not. Furthermore, from Equation 1, we know that an upper-bound on the total interfering workload suffered by task τ_{hi} is given by summing $\bar{I}_{j,hi}^{\text{NC}}(\Delta, \ell) \forall \tau_j \in \text{hp}(\tau_{hi})$ with the $m - 1$ largest values of $\bar{I}_{j,hi}^{\text{DIFF}}(\Delta, \ell)$. Thus, the actual total interfering workload suffered by τ_{hi} will be less than or equal to $\bar{I}_{hi}(\Delta, \ell)$. ■

Corollary 2. Upon a DMCR_ℓ , when the protocol identifies a job $J_{nh,s}$ that completed its execution at time $f_{nh,s} \leq r_{nh,s} + R_{nh}(\ell)$, then from time $f_{nh,s}$ onward, the actual total interfering workload suffered by any task $\tau_i \in \tau$ is less than or equal to $\bar{I}_i(\Delta, \ell)$.

From Corollary 2, we will now show that it is safe to re-enable every suspended task $\tau_k \in \tau^\ell$ at time $f_{nh,s}$.

Lemma 3. Upon a DMCR_ℓ , when protocol identifies a job $J_{nh,s}$ that completed its execution at time $f_{nh,s} \leq r_{nh,s} + R_{nh}(\ell)$, then the jobs released by every suspended task $\tau_k \in \tau^\ell$ from time $f_{nh,s}$ onward will meet their deadlines in any scenario of criticality level at most L_k .

Proof: The WCRT at criticality level ℓ of each task $\tau_k \in \tau^\ell$ was computed assuming a total interfering workload of $\bar{I}_i(\Delta, \ell)$. But from Corollary 2, we know that at time $f_{nh,s}$, the actual total interfering workload suffered by any task τ_i in the system is less than or equal to $\bar{I}_i(\Delta, \ell)$. Since the system was deemed MSM-schedulable, every job released by a task

$\tau_k \in \tau^\ell$ from time $f_{n_h,s}$ onward will meet its deadline in any scenario of criticality level at most L_k . ■

Lemma 4. Assume the system is executing in operating mode M_h . Upon a DMCR_ℓ , when protocol identifies a job $J_{n_h,s}$ that completed its execution at time $f_{n_h,s} \leq r_{n_h,s} + R_{n_h}(\ell)$, the re-enablement of every suspended task $\tau_k \in \tau^\ell$ will not jeopardize the h -periodicity and h -feasibility.

Proof: From Corollary 2, we know that at time $f_{n_h,s}$, the actual total interfering workload suffered by any task τ_i in the system is less than or equal to $\bar{I}_i(\Delta, \ell)$. However, this upper-bound is computed assuming that every task $\tau_k \in \tau^\ell$ will release jobs that complete their execution after no more than $C_i(\ell)$ time units. Therefore, the actual total interfering workload suffered by any task τ_j upon the re-enablement of every suspended task $\tau_k \in \tau^\ell$ will still be less than or equal to $\bar{I}_i(\Delta, \ell)$. Furthermore, since the system was deemed MSM-schedulable, it follows that every job released by tasks $\tau_j \in \tau^h$ will meet its deadline in any scenario of criticality up to level $L_j \geq h$. It follows that both the h -periodicity and h -feasibility are preserved upon the re-enablement of every suspended task $\tau_k \in \tau^\ell$. ■

Theorem 6. Let us assume the system is executing in operating mode M_h . Upon a DMCR_ℓ , when the protocol identifies a job $J_{n_h,s}$ that completed its execution at time $f_{n_h,s} \leq r_{n_h,s} + R_{n_h}(\ell)$, it is safe to re-enable the suspended task belonging to the operating mode M_ℓ at time $f_{n_h,k}$.

Proof: To prove this theorem, we have to show that Property 2 holds upon the re-enablement of every suspended task $\tau_j \in \tau^\ell$ at time $f_{n_h,k}$. Lemma 3 proved that every suspended task τ_i belonging to the operating mode M_ℓ could release jobs that would be able to meet their deadline if τ_i was re-enabled at time $f_{n_h,k}$, and provided that no $\text{IMCR}_{\ell+1}$ was triggered. Furthermore, Lemma 4 proved that the re-enablement of every suspended task τ_i belonging to the operating mode M_ℓ would not jeopardize the h -periodicity and h -feasibility. It follows that at time $f_{n_h,k}$, it is safe to re-enable every suspended task belonging to the operating mode M_ℓ . ■

Theorem 6 proves that it will eventually be possible to decrease the criticality level of the system, provided the computational demand of the tasks having a criticality higher than or equal to h decreases. It follows that the suspension delay suffered by tasks having a criticality less than h is reduced. In practice however, if an $\text{IMCR}_{\ell+1}$ is triggered during the DMCR_ℓ handling, then the procedure that consists in finding the first job $J_{n_h,s}$ that completes its execution a time $f_{n_h,s} \leq r_{n_h,s} + R_{n_h}(\ell)$ is aborted. Indeed, in that case, we can no longer guarantee that the additional interfering workload generated by tasks having a criticality equal to ℓ will not jeopardize the $(\ell + 1)$ -periodicity and $(\ell + 1)$ -feasibility.

Corollary 3. Upon a DMCR_ℓ , reenabling every suspended task belonging to operating mode M_ℓ upon the identification of a job $J_{n_h,s}$ that completed its execution at time $f_{n_h,s} \leq r_{n_h,s} + R_{n_h}(\ell)$ is a valid decreasing transition protocol.

IV. CONCLUSION

In this work, the first contribution consisted in formalizing mixed-criticality systems in terms of multi-moded systems, thus giving a new outlook to the problem. As a second contribution, we have shown that multi-moded approaches could help solve the consistency problems that arise when less critical tasks are brutally discarded, by enabling for a softer switch from a lower criticality level to a higher one. Finally, as a third contribution, we have highlighted the fact that the behavior of such systems could be greatly enhanced, by proving that task re-enablement was possible without compromising its safety. Those approaches allow to achieve a much more adept usage of the platform, by avoiding to vainly waste computational resources. Future work will concentrate, among other things, on suggesting asynchronous decreasing transition protocols, which allow for a progressive re-enablement of the less critical tasks.

REFERENCES

- [1] J. Real and A. Crespo, "Mode change protocols for real-time systems: A survey and a new proposal," *Real-Time Systems*, pp. 161–197, 2004.
- [2] S. Vestal, "Preemptive scheduling of multi-criticality systems with varying degrees of execution time assurance," in *RTSS 2007*, pp. 239–243.
- [3] H. Li and S. K. Baruah, "Load-based schedulability analysis of certifiable mixed-criticality systems," in *EMSOFT 2010*, pp. 99–108.
- [4] S. K. Baruah, V. Bonifaci, G. D'Angelo, A. Marchetti-Spaccamela, S. Van Der Ster, and L. Stougie, "Mixed-criticality scheduling of sporadic task systems," in *ESA 2011*, pp. 555–566.
- [5] S. K. Baruah, A. Burns, and R. Davis, "Response-time analysis for mixed criticality systems," in *RTSS 2011*, pp. 34–43.
- [6] N. Guan, P. Ekberg, M. Stigge, and W. Yi, "Effective and efficient scheduling of certifiable mixed-criticality sporadic task systems," in *RTSS 2011*, pp. 13–23.
- [7] S. K. Baruah, V. Bonifaci, G. D'Angelo, H. Li, A. Marchetti-Spaccamela, S. van der Ster, and L. Stougie, "The preemptive uniprocessor scheduling of mixed-criticality implicit-deadline sporadic task systems," in *ECRTS 2012*, pp. 145–154.
- [8] F. Santy, L. George, P. Thierry, and J. Goossens, "Relaxing mixed-criticality scheduling strictness for task sets scheduled with FP," in *ECRTS 2012*, pp. 155–165.
- [9] R. M. Pathan, "Schedulability analysis of mixed-criticality systems on multiprocessors," in *ECRTS 2012*, pp. 309–320.
- [10] H. Li and S. K. Baruah, "Global mixed-criticality scheduling on multiprocessors," in *ECRTS 2012*, pp. 166–175.
- [11] N. Audsley, A. Burns, M. Richardson, K. Tindell, and A. J. Wellings, "Applying new scheduling theory to static priority pre-emptive scheduling," *Software Engineering Journal*, vol. 8, pp. 284–292, 1993.
- [12] N. Guan, M. Stigge, W. Yi, and G. Yu, "New response time bounds for fixed priority multiprocessor scheduling," in *RTSS 2009*, pp. 387–397.
- [13] S. K. Baruah, V. Bonifaci, G. D'Angelo, H. Li, A. Marchetti-Spaccamela, N. Megow, and L. Stougie, "Scheduling real-time mixed-criticality jobs," *IEEE Transaction on Computers*, pp. 1140–1152, 2012.
- [14] S. K. Baruah, H. Li, and L. Stougie, "Mixed-criticality scheduling: improved resource-augmentation results," in *CATA 2010*, pp. 217–223.
- [15] —, "Towards the design of certifiable mixed-criticality systems," in *RTAS 2010*, pp. 13–22.
- [16] H. Li and S. K. Baruah, "An algorithm for scheduling certifiable mixed-criticality sporadic task systems," in *RTSS 2010*, pp. 183–192.
- [17] P. Ekberg and W. Yi, "Bounding and shaping the demand of mixed-criticality sporadic tasks," in *ECRTS 2012*, pp. 135–144.
- [18] V. Nelis, J. Goossens, and B. Andersson, "Two protocols for scheduling multi-mode real-time systems upon identical multiprocessor platforms," in *ECRTS 2009*, pp. 151–160.
- [19] L. Cucu-Grosjean and J. Goossens, "Predictability of fixed-job priority schedulers on heterogeneous multiprocessor real-time systems," *Information Processing Letters*, vol. 110, no. 10, pp. 399–402, 2010.